

WHAT IS CLAIMED IS:

1. A method for solving, in a computer, stochastic control problems of linear systems in high dimensions, comprising:

(a) modeling, in the computer, a structured Markov Decision Process (MDP),
5 wherein a state space for the MDP is a polyhedron in a Euclidean space and one or more actions that are feasible in a state of the state space are linearly constrained with respect to the state; and

(b) building, in the computer, one or more approximations from above and from below to a value function for the state using representations that facilitate the
10 computation of approximately optimal actions at any given state by linear programming.

2. The method of claim 1, wherein the MDP comprises a supply chain planning process.

15 3. The method of claim 1, wherein the action space and the state space are continuous and related to each other through a system of linear constraints.

4. The method of claim 1, wherein the value function is convex and the method further comprises efficiently learning the value function in advance and
20 representing the value function in a way that allows for real-time choice of actions based thereon.

5. The method of claim 1, wherein the linear function is approximated both from above and from below by piecewise linear and convex functions.

6. The method of claim 5, wherein the domains of linearity of the piecewise linear and convex functions are not stored explicitly, but rather are encoded through a linear programming formulation.

7. The method of claim 5, wherein the domains of linearity of the piecewise linear and convex functions allow the functions to be optimized and updated in real-time.

8. The method of claim 1, wherein the value function can be efficiently approximated both from above and from below.

9. The method of claim 1, wherein the approximations can be repeatedly refined.

10. The method of claim 1, wherein the value function can be efficiently approximated from above based on knowledge of upper bounds on the function at each member of a selected set of states.

11. The method of claim 1, wherein the value function can be efficiently approximated from below based on linear functions that lie below the convex value function.

5 12. The method of claim 1, wherein the value function can be approximated successively.

13. A computerized apparatus for solving stochastic control problems of linear systems in high dimensions, comprising:

10 (a) a computer;

(b) logic, performed by the computer, for modeling a structured Markov Decision Process (MDP), wherein a state space for the MDP is a polyhedron in a Euclidean space and one or more actions that are feasible in a state of the state space are linearly constrained with respect to the state; and

15 (c) logic, performed by the computer, for building one or more approximations from above and from below to a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming.

20 14. The apparatus of claim 13, wherein the MDP comprises a supply chain planning process.

15. The apparatus of claim 13, wherein the action space and the state space are continuous and related to each other through a system of linear constraints.

16. The apparatus of claim 13, wherein the value function is convex and the
5 logic further comprises efficiently learning the value function in advance and representing the value function in a way that allows for real-time choice of actions based thereon.

17. The apparatus of claim 13, wherein the linear function is approximated
10 both from above and from below by piecewise linear and convex functions.

18. The apparatus of claim 17, wherein the domains of linearity of the
piecewise linear and convex functions are not stored explicitly, but rather are encoded
through a linear programming formulation.

19. The apparatus of claim 17, wherein the domains of linearity of the
piecewise linear and convex functions allow the functions to be optimized and updated in
real-time.

20. The apparatus of claim 13, wherein the value function can be efficiently
20 approximated both from above and from below.

21. The apparatus of claim 13, wherein the approximations can be repeatedly refined.

22. The apparatus of claim 13, wherein the value function can be efficiently approximated from above based on knowledge of upper bounds on the function at each member of a selected set of states.

23. The apparatus of claim 13, wherein the value function can be efficiently approximated from below based on linear functions that lie below the convex value function.

24. The apparatus of claim 13, wherein the value function can be approximated successively.

25. An article of manufacture embodying logic for solving stochastic control problems of linear systems in high dimensions, the logic comprising:

- (a) modeling a structured Markov Decision Process (MDP), wherein a state space for the MDP is a polyhedron in a Euclidean space and one or more actions that are feasible in a state of the state space are linearly constrained with respect to the state; and
- (b) building one or more approximations from above and from below to a value function for the state using representations that facilitate the computation of approximately optimal actions at any given state by linear programming.

26. The article of manufacture of claim 25, wherein the MDP comprises a supply chain planning process.

5 27. The article of manufacture of claim 25, wherein the action space and the state space are continuous and related to each other through a system of linear constraints.

28. The article of manufacture of claim 25, wherein the value function is
10 convex and the logic further comprises efficiently learning the value function in advance and representing the value function in a way that allows for real-time choice of actions based thereon.

29. The article of manufacture of claim 25, wherein the linear function is
15 approximated both from above and from below by piecewise linear and convex functions.

30. The article of manufacture of claim 29, wherein the domains of linearity of the piecewise linear and convex functions are not stored explicitly, but rather are encoded through a linear programming formulation.

20

31. The article of manufacture of claim 29, wherein the domains of linearity of the piecewise linear and convex functions allow the functions to be optimized and updated in real-time.

5 32. The article of manufacture of claim 25, wherein the value function can be efficiently approximated both from above and from below.

10 33. The article of manufacture of claim 25, wherein the approximations can be repeatedly refined.

34. The article of manufacture of claim 25, wherein the value function can be efficiently approximated from above based on knowledge of upper bounds on the function at each member of a selected set of states.

15 35. The article of manufacture of claim 25, wherein the value function can be efficiently approximated from below based on linear functions that lie below the convex value function.

20 36. The article of manufacture of claim 25, wherein the value function can be approximated successively.